Can one Patient Record Accommodate the Diversity of Specialized Care?

Astrid M. van Ginneken, M.D. Ph.D., Huibert Stam, B.Sc.
Department of Medical Informatics
Erasmus University, Rotterdam, The Netherlands

Despite a quarter century of developments, few specialists directly use a computerized patient record, that fully replaces the paper chart. Because of the diversity of domains in specialized care, medical decision-making and the continuity of care may suffer from scattering of patient data over various records. The challenge was to develop a computerized patient record, that would be versatile enough to tailor it to specific needs, while keeping it uniform enough to permit physicians to share data on the same patient. In our CPR, the key that reconciles versatility with uniformity lies in the design of the data model. The CPR consists of a mother record with specialized sub-records, that all share the same data model. A physician can enlarge his scope for decision-making by consulting other specialized records on the same patient or by viewing the combined information of all sub-records without the need to convert data or to familiarize himself with different interfaces.

INTRODUCTION

The growing body of medical knowledge and the development of new techniques for diagnosis and treatment, have led to (super)-specialization and an increasing complexity of health care. As a result, the shortcomings of the traditional paper medical record (PMR) are rapidly outweighing its strengths [1]. Advantages of PMRs like portability, flexibility, and browsability are turning into limitations. The paper record can be moved around, but it can only be in one location. Flexibility makes retrieval and analysis cumbersome, because it gives way to poor organization and incompleteness. Browsing looses its effectiveness when dealing with illegible handwriting of colleagues and scattering of patient data over several volumes at different departments. Advances in computer technology have inspired researchers for a quarter century to create a versatile record for the benefit of patient care and research. Although the introduction of computer-based patient records (CPRs) in primary care has been successful in some areas [2], there is no CPR in specialized care, that is used interactively by physicians and that fully replaces the paper chart. Electronic storage of textual data does provide a solution to legibility and availability, but structure is required to support versatile views on the data, decision-support, and data analysis. Developers of medical records are faced with the tension between effort and benefit when it comes to the collection of structured data from physicians. Which strategy is most effective in supporting data entry depends on the clinical setting and the size of the medical domain [3]. The crux of the problem is to confront physicians with data sets, that are tailored to the tasks they have to perform [4]. Data entry is not the only problem. In specialized care there are often as many PMRs as there are specialties. The resulting scattering of data hinders treating-physicians in the formation of a complete picture of the patient. The continuity and quality of care will benefit from the possibility of shared data in the records of different specialists. Hence, in a field as diverse as specialized care, it is not desirable to develop separate equivalents of the present PMRs. The challenge is to create a CPR with as many 'faces' as there are specialties. This paper describes and illustrates a data model underlying such a versatile record.

THE DATA MODEL

Considerations for the design of the data model

The design of the data model was preceded by the formulation of the functional requirements of the CPR. Issues related to temporal aspects, the interface of the CPR, and support of data entry have been described elsewhere [3, 5] and will not be discussed in this paper. We will only focus on those requirements and aspects of the design that are related to the versatility needed in a specialized care setting.

Efficient data entry. The physician should not be forced, but stimulated to enter data in a structured fashion by drawing attention to the benefits of structure[5]. Therefore, data entry must be supported at several levels of detail. Views for data entry must conform to sets of data, as collected in clinical practice.

Consultation. The CPR should provide versatile views that are tailored to the medical setting - such as first visit, follow-up visit, and function test - and the problem-solving task at hand [4]. In addition to predefined views, the CPR should support flexible browsing, like zooming in on certain types of information and searching back and forth in time.

Sharing of records. The structure of the CPR must support the option to consult other CPRs of the same patient without the need for data conversions and familiarization with different interfaces. There should also be the potential to combine data of CPRs of different specialists into one view. This *patient view* crosses the boundary of one specialty and informs the physician about all diagnoses, medications, and test results, applicable to a particular patient.
These requirements are seemingly in conflict. How can the diversity of specialized care and the necessity to tailor CPR applications to clinical practice be reconciled with the uniformity required for sharing interface and contents? This reconciliation can only be found in a model that can accommodate a variety in content with a uniform structure. The variety in content offers the flexibility to tailor CPR applications to the requirements of different specialties, whereas a uniform structure permits transparent views on the data of one patient in CPRs of different specialists.

Design of the data model

Earlier, we described the idea of a mother record, which is extendable with sub-records, each tailored to a specific domain [5]. The mother record is used to access the sub-records and allows specialists to record information, that is not covered by their specialized sub-record. Both the mother record and the sub-records share the same data model. The main building blocks of the data model are the *event* and the *action*. The event denotes a set of data, that belongs together in the sense that it fulfills three requirements: entered at one moment, discovered at one moment, and originating from one source. A typical example of an event is the patient visit, but the arrival of a lab result, or pathology report, is also an event. An event may contain more than one action, depending on the type of event. The event patient visit will usually contain actions such as history, physical exam, and medication. Events and actions are used to represent the facts that stem from what the physician observed, thought, and did. These facts correspond to level 1 in Rector's model [6].

In order to store level 1 patient data, we developed a relational model, using the Boyce/Codd normal form [7]. Figure 1 shows an ERD (entity relationship diagram) of the essence of the data model.

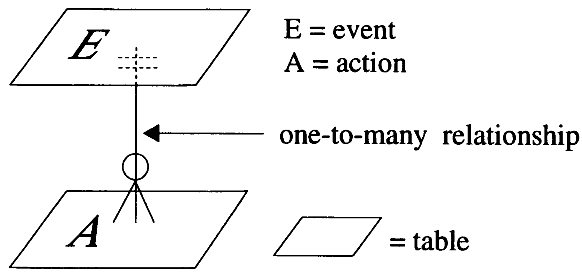

Figure 1. The core structure of the data model. Each event can be composed of one or more actions.

The potential to define specialized data sets is achieved by defining subtypes of the supertypes eventand action. In this way, one can create super-subclass trees, which resemble hierarchies of classes in object-oriented models. The extensibility of the model lies in the creation of new subtypes of events and actions. Figure 2 shows the ERD of Figure 1, extended with several subtypes. We will refer to these subtypes as sub-events and sub-actions. Tables with a super-subtype connection have the following properties:

- Each record in a subtype table corresponds with exactly one record (same key attributes) in all its supertype tables
- Each record in a supertype table has a corresponding record (same key attributes) in exactly one of its direct subtype tables

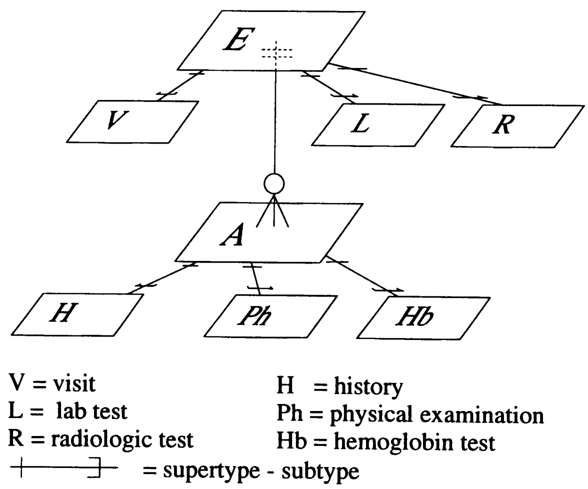

Figure 2. The supertypes 'event' and 'action' each serve as the top node of their own hierarchy of specialized subtypes.

Taking these properties into account, it can be seen that the database structure, as depicted in Figure 2, would allow each type of action to belong to each type of event. Although this is extremely flexible, there are many combinations that do not make sense in clinical practice: an action representing a chest X-ray report should not be part of a patient visit. Similarly, the prescription of a certain drug should not be part of a lab result.
It is possible to define constraints, that determine which sub-actions may belong to which sub-events. However, when these constraints are explicitly defined at the level of the data model, i.e. as database constraints, then an explosion of tables and redundancy would be the result: there would have to be as many tables of a sub-action as there are sub-events to which that sub-action may belong. Such a data model would soon become huge and difficult to maintain. For the sake of simplicity, flexibility, and maintenance, we have chosen to define the necessary constraints at the level of the interface. Within a screen that corresponds with a certain type of event, the user will only have access to those sub-actions that may belong to that event. This strategy of event and action hierarchies, combined with constraints at the interface level produces an extremely flexible instrument to define specialized data sets and combine them into dedicated records.

Application of the data model
At present, a mother record has been created and two specialized records in the domains of heartfailure and andrology [8]. Figure 3 shows a simplified ERD of the data model and the interface constraints.

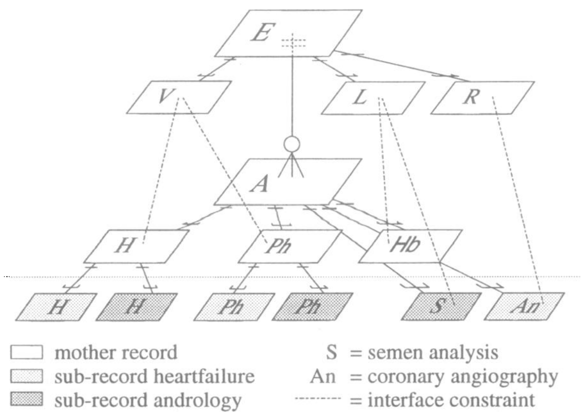

Figure 3. The actions of the mother record have been extended with subtypes for specialized records. The interface constraints realize a proper linkage between events and actions.

From this picture, the question arises, which mechanism determines how subtypes of events and actions are combined to form specialized records. The andrologist will want access to the andrology-specific history, whereas the cardiologist needs the heartfailure-specific history. Both specialized histories belong to the event 'patient visit'. The task to navigate through the mother record and the appropriate sub-record is performed by the 'state transition manager' (STM). Each screen for consultation or data entry represents a state in the interface and a view to one or more of the sub-action tables in the database. The STM contains knowledge about navigation through the various screens. In other words, when the user leaves a screen, the STM consults the state-transition table for the next state, i.e. screen, to be selected. The state-transition table can easily be extended as new subtypes and screens are added to the application. When the user proceeds from the mother record to one of the sub-records, the number of potential transitions depends on the number of sub-records. By default, the STM will use the identity of the user and the selected patient to select the appropriate sub-record. With this mechanism, the mother record will automatically connect with the subtypes of the specialized record that the physician most often needs. When zooming in on history from the 'patient visit' screen in the mother record, the andrologist will proceed with the andrology history screen and the cardiologist with the heartfailure history screen: this is schematically illustrated in Figure 4. If a different sub-record needs to be accessed, the physician has to specify which one from a menu, while still at the level of the mother record. If the physician is confronted with information outside his own domain of expertise, he can record it at the level of the mother record.

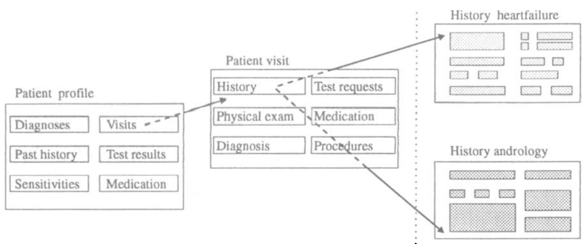

Figure 4. From the patient profile screen, the physician can zoom in to see more detailed data. The history screens of two specialized sub-records are to the left of the dotted line. The mother record is to the right. The application selects the appropriate screens when the physician proceeds to the specialized level.

The data model supports the option to share information with other specialists. Provided that physicians give read-only access to others who treat the same patient, it is possible to access the specialized records of other specialists to form a more complete picture of that patient. There is also the option for a more direct sharing of information on a particular patient. As opposed to the 'normal' mode in which a physician only sees information that he recorded himself, the user can switch to 'patient view' mode at the level of the mother record. The scope of this view includes the events of all existing sub-records on that patient. At the database level this is no more than removing the constraint that the data to be displayed was recorded by the physician who logged in. In patient view mode, the overview screen 'patient profile' will show and provide access to all diagnoses, visit dates, medications, and tests results of all physicians treating that patient.

DISCUSSION

The objective was to develop a CPR that would allow for efficient data entry, consultation, and sharing of information in a specialized care setting. Hence, we were faced with the challenge to create a CPR that would be versatile enough to tailor it to the specific needs of different specialists while keeping it uniform enough for sharing interface and contents.

Other researchers have described models for the representation of patient data. In his framework for the patient record, Rector [9] introduces 'individuals' which are observable entities, such as 'Jane Smith's fracture'. 'Occurrences' are observations on an individual by an agent at a particular place and time. 'Categories' are concepts from which individuals form instantiations: Jane Smith's fracture is an instantiation of the general concept 'fracture'. Moorman [10] worked on modelling of description knowledge, that defines how where and when concepts can be described. Description knowledge is intended to support data entry. The actual patient data are stored in a hierarchy of instantiated concepts where the paths in the hierarchy represent the findings. Dolin [11] proposed the use of nested triples to describe complex findings, which are composed of more than one observation. Bernauer [12] and Campbell [13] use conceptual graphs to define how concepts can be combined to represent potential medical findings. The findings themselves are instantiations of these graphs. All these models allow for various levels of detail when recording patient data and they can all be applied to different domains.

Yet, these models are difficult to compare with our data model, because they were described with a different focus. The publications do not explicitly reveal how data model and interface are related and how they are used to create and combine data items into views that different specialists can identify with and appreciate.

Several applications, such as Ivory [14], Pen&Pad [15], IMR-E [16], and Pure-MD [17] have dynamic interfaces, in the sense that the options for data entry depend on previous selections by the physician. The knowledge underlying the interface represents findings that may be described in one or more medical domains. The fact that any domain can be modelled at a high level of detail does not automatically imply that medical concepts can be combined into different views to accommodate the working styles of different specialists. However, Ivory offers a mechanism to tailor its interface to the needs of a specialist or researcher: the user can define a view in which he combines the medical concepts that he needs to describe within a certain context. That context is defined as a problem and the view is invoked when that problem is assigned to a patient. A similar mechanism is realized in Pure-MD through the use of pre-defined encounter filters. Yet, it is unclear whether or not these applications permit physicians to consult each other's records or support views on a patient's data beyond the scope of one specialty.

The key to the versatility and uniformity of our CPR for specialized care lies in the design of our data model. We have shown that the CPR can be tailored to specific needs by defining specializations of categories of information, combined with interface constraints at screen level, and knowledge for navigation through the screens. This mechanism of creating hierarchies of specializations, offers the potential to record and consult data at levels of increasing detail within a specialized view.
Browsing records of other specialists and sharing information of different specialists in one view is easy because of the general structure of the data model. Specialized sub-records only differ in the contents of their screens and the attributes of their corresponding actions. The mother record and the navigation through the screens is the same for all dedicated parts of the CPR. Therefore, physicians have no need to familiarize themselves with different interfaces. The option of a 'patient view' offers a new perspective for decision-support. For example, detection of interactions and contra-indications for drugs no longer need to be restricted to data within one specialized

record, but can include diagnoses, test results, and prescriptions of co-treating specialists.
The CPR can further be tailored to specific needs by exploiting its multi-media potential when embedded in a medical workstation, like HERMES [18]. HERMES is a client-server architecture that can access other data-sources, such as echo, angio, and ECG-collectors, and present these images and signals as part of the CPR. Direct availability of non-textual sources of information on a patient will reduce the time needed for physicians to decide on further investigations and treatment. Which functionality of the CPR will be available to physicians greatly depends on the infrastructure and cooperation of the institution in which the CPR is used.

In conclusion, exploitation of our data model may ultimately improve the continuity of care by sharing of information among specialists and contribute to the quality of care by a wider scope for decision-making.

References

1. Dick RS, Steen EB. The computer-based patient record: an essential technology for health care. Committee on Improving the Patient record Division of Health Care Services. Institute of medicine. National Academy Press 1991.
2. van der Lei J, Duisterhout JS, Westerhof HP, et al. The introduction of computer-based patient records in The Netherlands [see comments]. Ann Intern Med 1993;119(10):1036-41.
3. van Ginneken AM, Stam H, Moorman PW. A multi-strategy approach for medical records of specialists. In: Greenes RA, Peterson HE, Protti DJ. Proceedings of MEDINFO 95, Vancouver B.C. 1995:290-4.
4. Nygren E, Henriksson P. Reading the medical record. I. Analysis of physicians' ways of reading the medical record. Comput Methods Programs Biomed 1992;39(1-2):1-12.
5. van Ginneken AM, Stam H, Duisterhout JS. A powerful macro-model for the computer patient record. Proc Annu Symp Comput Appl Med Care 1994:496-500.
6. Rector AL, Nowlan WA, Kay S. Foundations for an electronic medical record. Methods Inf Med 1991;30(3):179-86.
7. Date CJ. An Introduction to Database Systems. (4th ed.) Reading, MA: Addison-Wesley, 1986. (Howe ME, ed. The Systems Programming Series; vol 1).
8. Timmers T, Pierik F, Steenbergen M, Stam H, van Ginneken AM, van Mulligen EM. ARIS: Integrating Multi-source Data for Research in Andrology. Accepted for the 19th SCAMC.
9. Rector AL, Nowlan WA, Kay S, Goble CA, Howkins TJ. A framework for modelling the electronic medical record. Methods Inf Med 1993;32(2):109-19.
10. Moorman PW, van Ginneken AM, van der Lei J, van Bemmel JH. A model for structured data entry based on explicit descriptional knowledge. Methods Inf Med 1994;33(5):454-463.
11. Dolin RH. Modeling in relational complexities of symptoms. Methods Inf Med 1994;33(5):448-453.
12. Bernauer J. Conceptual graphs as an operational model for descriptive findings. In: Clayton PD, ed. Proceedings of the 15th SCAMC. New York: McGraw-Hill 1991:214-218.
13. Campbell KE, Das AK, Musen MA. A logical foundation for representation of clinical data. J Am Med Informatics Assoc 1994;1(3):218-232.
14. Campbell KE, Musen MA. Creation of a systematic domain for medical care: the need for a comprehensive patient-description vocabulary. In: Lun KC et al. (eds). Proceedings of MEDINFO92 1992:1437-1442.
15. Nowlan WA, Rector AL, Kay S, Horan B, Wilson A. A patient care workstation based on user centred design and a formal theory of medical terminology: PEN&PAD and the SMK formalism. Proc Annu Symp Comput Appl Med Care 1991:855-7.
16. Trace D, Naeymi-Rad F, Haines D, et al. Intelligent Medical Record--entry (IMR-E). J Med Syst 1993;17(3-4):139-51.
17. Lussier YA, Maksud M, Desruisseaux B, Yale PP, St-Arneault R. PureMD: a Computerized Patient Record software for direct data entry by physicians using a keyboard-free pen-based portable computer. Proc Annu Symp Comput Appl Med Care 1992:261-4.
18. van Mulligen EM, Timmers T, van Bemmel JH. A new architecture for integration of heterogeneous software components. Methods Inf Med 1993;32(4):292-301.

Address of correspondence:

Astrid M. van Ginneken M.D. Ph.D.
Department of medical Informatics
Erasmus University
P.O. Box 1738, 3000 DR Rotterdam
The Netherlands
E-mail: vanginneken@mi.fgg.eur.nl